920476-904898

## IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

| | | |
|---|---|---|
| **In the application of** | : | Weil et al |
| **Serial No.** | : | 09/848,743 |
| **Filed** | : | May 3, 2001 |
| **For** | : | Communication System and Method of Optimising Provisioning of Service to Subscribers in a Passive Optical Network |
| **Examiner** | : | Salad, E. Abdullahi |
| **Art Unit** | : | 2157 |
| **Customer number** | : | 23644 |

## SUBMISSION OF AMENDED BRIEF ON APPEAL

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

In accordance with the Notice mailed October 16, 2006, submitted herewith is an Amended Appeal Brief, having appropriate references to the specification added to the Summary Section of the Brief.

November 13, 2006

Respectfully submitted,

William M. Lee

William M. Lee, Jr.
Registration No. 26,935
Barnes & Thornburg
P.O. Box 2786
Chicago, Illinois 60690-2786
(312) 214-4800
Fax (312) 759-5646

920476-904898

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

| | | |
|---|---|---|
| **In the application of** | : | Weil et al |
| **Serial No.** | : | 09/848,743 |
| **Filed** | : | May 3, 2001 |
| **For** | : | Communication System and Method of Optimising Provisioning of Service to Subscribers in a Passive Optical Network |
| **Examiner** | : | Salad, E. Abdullahi |
| **Art Unit** | : | 2157 |
| **Customer number** | : | 23644 |

**AMENDED BRIEF ON APPEAL**

Honorable Director of Patents and Trademarks
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

This appeal is from the Examiner's final rejection of November 2, 2005 and the advisory action of February 1, 2006 in which all pending claims, that is claims 1-28 were rejected or objected to. An appropriate response was mailed on January 3, 2006, and a timely Notice of Appeal was mailed on March 2, 2006 with the required fee of $500.

The Brief fee of $500 was deducted from Deposit Account No. 12-0913 when the brief was originally filed October 6, 2006. A petition for a two month extension of time accompanied the earlier-filed Brief.

## (i) Real Party in Interest

This application is assigned to Nortel Networks Limited.

## (ii) Related Appeals and Interferences

There are no related appeals or interferences or judicial proceedings.

## (iii) Status of Claims

This application was filed with claims 1 through 28. Claims 1-28 have been finally rejected by the Examiner. The rejection of these claims in the office action of November 2, 2005 is appealed. Claims 1-28 are set forth in the Claims Appendix.

## (iv) Status of Amendments

Amendments were made in the after final response of January 3, 2006, and were entered.

## (v) Summary of Invention

The present invention relates generally to multi layer communication networks, and more specifically to methods of managing such networks involving identifying a protection path so as to remedy a fault condition. The communications network can have a lower transport layer and an upper overlay incorporating a plurality of routers. Adjacencies are defined between respective pairs of routers. Protection paths are identified by defining a software model of the overlay of the network and defining in the model a hierarchy of protection levels for the main path. Each protection level is characterised by a respective set of one or more broken adjacencies in the model,

the broken adjacencies meaning assumed failures of certain network elements. Which network elements being assumed to be broken is determined from a knowledge of the underlying topology. A protection level is selected and a protection path avoiding the broken adjacencies associated with that protection level is calculated. If this calculated path is not available in the real network, a lower protection level in the hierarchy is selected and the protection path calculation repeated. The process continues until a suitable protection path is identified.

<u>Independent claim 1</u>:

This claim specifies a method of fault recovery in a multi-layer communications network (figs 1a, 3a-3c, lines 17-25 of page 6) having a transport layer topology (fig 3b, lines 20-25 of page 8) and an overlay topology (fig 3a, lines 17-20 of page 8), in which adjacencies (12, 14, lines 22-25 of page 8) are defined between a plurality of network nodes, the method comprising, for each adjacency for which a recovery path is to be determined, modifying (fig 11, 102, "level of assumption" 105, "edit overlay topology" line 30 of page 14 onwards) the overlay topology by removal of selected adjacencies, attempting computation (106 "compute recovery path" line 7 of page 15 onwards) of a path, and if no path is available (107 "recovery path found? no", line 10-20 page 15) removing fewer selected adjacencies (108 "level of assumption refined", 105"edit overlay topology", line 11-16 page 15) from the overlay topology and repeating said path computation (106 "compute recovery path", lines 13-16 of page 15, line 7 onwards of page 15).

References to items shown in figures 1a, 2a, 3a, 3b and 11 have been added to show where examples of these claim features are illustrated. Other figures are referred to as showing more details such as examples of protection levels having fewer selected adjacencies removed.

Figures 3a to 3c illustrate the steps in the construction of the model of figure 2a from the real network topology. Figure 3a depicts an overlay topology in the form of a Layer 3 topology of the network and shows the Layer 3 interconnections 14 between the routers disposed at the network nodes A to L. Figure 3b depicts a transport layer

topology in the form of an underlying Layer 1 topology and shows the physical interconnections 12 between the switches disposed at the network nodes. These two topologies are combined in the composite diagram of figure 3c in which the Layer 3 interconnections or adjacencies 14 are shown explicitly, and the Layer 1 interconnections are inferred implicitly from physical proximity between nodes.

Fig 11 shows a flow chart of steps in calculating protection paths. With knowledge of the physical transport and the logical overlay topologies, it starts from an assumption of any logical path failure, and computes all the possible physical causes of that failure, then the method removes from the logical topology all possibly affected logical components and computes a recovery logical path which avoids all possible fault locations (which implies an optimum level of protection). If, as is frequently the case, this first choice is not available, then fewer of the possibly affected logical components are removed from the logical overlay topology, and another recovery logical path is determined which may use a possibly faulty physical element (which implies a weaker level of protection). Hence the method involves computing recovery paths successively in order of strength from full to weaker ones. This can enable more efficient computation compared to computing many paths then trying to determine which are stronger or weaker.

Independent claim 3

This claim specifies a method of calculating a protection path for traffic carried on a main path in a multilayer communications network (figs 1a, 3a-3c, lines 17-25 of page 6) having a lower transport layer and an upper layer incorporating a plurality of routers (13A-13D, lines 8-14 of page 7, lines 20-25 of page 8), and in which adjacencies (12) are defined between respective pairs of routers, the method comprising the steps of;

defining a model (figs 3a-3c, lines 17-25 of page 8) of the network;

defining in said model a hierarchy of protection levels (fig 2b, lines 1-15 of page 8), each said protection level being characterised by a respective set of broken adjacencies in said model; (lines 10-23 of page 10) attempting to calculate (fig 11,

106, lines 7-16 of page 15) a recovery path for a selected protection level in said hierarchy; and

if no said path is available (107, fig 11, lines 7-16 of page 15), repeating said calculation attempt (108, 106, fig 11, lines 7-16 of page 15) for successive further protection levels in said hierarchy until a protection path is identified.

References to items shown in figures 1a, 2b, 3a, 3b, 3c and 11 have been added to show where examples of these claim features are illustrated. Similar comments to those for claim 1 apply.

<u>Independent claim 4</u>

This claim specifies a method of calculating a protection path for traffic carried on a main path in a multi-layer communications network (Figs 1a, 3a-3c, lines 17-25 of page 6) comprising a lower transport layer and an upper overlay incorporating a plurality of routers (13A-13D, lines 8-14 of page 7, lines 20-25 of page 8), there being a plurality of adjacencies (12) defined between respective pairs of routers, wherein the method comprises the steps of;

defining a software model (Fig 3a, lines 17-25 of page 8) of the overlay of said network;

defining in said model a hierarchy of protection levels (Fig 2b, lines 1-15 of page 8) for said main path, each said protection level being characterised by a respective set of one or more broken adjacencies in said model;

selecting one said protection level and calculating (106, fig 11, lines 7-16 of page 15) a protection path avoiding the broken adjacencies associated with that protection level;

determining (107, fig 11, lines 7-16 of page 15) whether the calculated protection path is available in the network; and,

if said calculated path is not available in the network, repeating (108, 106, fig 11, lines 7-16 of page 15) said path calculation (106) and determining (107) steps for one or more further selected protection levels.

References to items shown in figures 1a, 2b, 3a, 3b, 3c and 11 have been added to show where examples of these claim features are illustrated. Similar comments to those for claim 1 apply.

<u>Independent claim 15</u>

This claim specifies a method of identifying a protection path for traffic carried on a main path in a multi-layer communications network (Figs 1a, 3a-3c) so as to remedy a fault condition involving one or more network elements (13A-13D), the method comprising the steps of:

assuming the failure (102, fig 11, line 30 of page 14 onwards) of a plurality of network elements;

calculating (106, fig 11, lines 7-16 of page 15) a protection path circumventing all the assumed failed network elements and determining (107, fig 11, lines 7-16 of page 15) whether that calculated path is an available path in the network; and

if said calculated path is not an available path in the network, successively assuming the failure of fewer network elements (108, fig 11, lines 7-16 of page 15) and repeating the path calculation (106) and determination (107) steps for each successive assumption until an available protection path is identified.

Again references to items shown in figures 1a, 2b, 3a, 3b, 3c and 11 have been added to show where examples of these claim features are illustrated. Similar comments to those for claim 1 apply.

<u>Independent claim 16</u>

This claim specifies software in machine readable form for identifying a protection path for traffic carried on a main path in a multi-layer communications network (figs 1a, 3a-3c, lines 17-25 of page 6) so as to remedy a fault condition involving one or more network elements (13A-13D, lines 8-14 of page 7, lines 20-25 of page 8), the software being arranged to perform the method steps of:

assuming (102) the failure of a plurality of network elements;

calculating (106, fig 11, lines 7-16 of page 15) a protection path circumventing all the assumed failed network elements and determining (107, fig 11, lines 7-16 of page 15) whether that calculated path is an available path in the network; and

if said calculated path is not an available path in the network, successively assuming (108, fig 11, lines 7-16 of page 15) the failure of fewer network elements and repeating the path calculation (106) and determination (107) steps for each successive assumption until an available protection path is identified.

Again, references to items shown in figures 1a, 2b, 3a, 3b, 3c and 11 have been added to show where examples of these claim features are illustrated. Similar comments to those for claim 1 apply.

<u>Independent claim 17</u>

This claim specifies a method of network management for planning fault recovery paths in a multi-layer communications network (Figs 1a, 3a-3c, lines 17-25 of page 6) having a transport layer (Fig 3b, lines 20-25 of page 8) and an overlay topology (fig 3a, lines 17-20 of page 8), and in which adjacencies (12,14, lines 22-25 of page 8) are defined between a plurality of network nodes, the method comprising, providing a model (Fig 3c, lines 17-25 of page 8) of the network, and within said model for each recovery path to be determined, modifying (102, fig 11, line 30 of page 14 onwards) the overlay topology by removal of selected adjacencies, attempting computation (106, fig 11, lines 7-16 of page 15) of a recovery path, and if no recovery path is available (107, fig 11, lines 7-16 of page 15) removing fewer selected adjacencies (108) from the overlay topology and repeating (106, fig 11, lines 7-16 of page 15) said recovery path computation.

Again, references to items shown in figures 1a, 2b, 3a, 3b, 3c and 11 have been added to show where examples of these claim features are illustrated. Similar comments to those for claim 1 apply.

Independent claim 18

This claim specifies a method of network management for planning fault recovery paths in a multi-layer communications network (Fig 1a, lines 17-25 of page 6) so as to remedy a future fault condition involving one or more network elements (13A-13D, lines 8-14 of page 7, lines 20-25 of page 8) on a traffic path, the method comprising the steps of:

providing a model (figs 3a-3c, lines 17-25 of page 8) of the network

assuming (102, fig 11, line 30 of page 14 onwards) the failure of a plurality of network elements in said model;

calculating (106, fig 11, lines 7-16 of page 15) a protection path circumventing all the assumed failed network elements and determining (107, fig 11, lines 7-16 of page 15) whether that calculated path is an available path in the network; and

if said calculated path is not an available path in the network, successively assuming (108, fig 11, lines 7-16 of page 15) the failure of fewer network elements and repeating the path calculation (106) and determination (107) steps for each successive assumption until an available protection path is identified.

Again, references to items shown in figures 1a, 2b, 3a, 3b, 3c and 11 have been added to show where examples of these claim features are illustrated. Similar comments to those for claim 1 apply.

Independent claim 19

This claim specifies a network manager for a multi-layer communications network (figs 1a, 3a-3c, lines 17-25 of page 6) and for planning fault recovery paths in said network so as to remedy a fault condition involving one or more network elements (13A-13D, lines 8-14 of page 7, lines 20-25 of page 8) on a traffic path, the network manager being arranged to perform the method steps of:

providing a model (figs 3a-3c, lines 17-25 of page 8) of the network

assuming (102, fig 11, line 30 of page 14 onwards) the failure of a plurality of network elements in said model;

calculating (106, fig 11, lines 7-16 of page 15) a protection path circumventing all the assumed failed network elements and determining (107, fig 11, lines 7-16 of page 15) whether that calculated path is an available path in the network; and
if said calculated path is not an available path in the network, successively assuming (108, fig 11, lines 7-16 of page 15) the failure of fewer network elements and repeating the path calculation (106) and determination (107) steps for each successive assumption until an available protection path is identified.
Again, references to items shown in figures 1a, 2b, 3a, 3b, 3c and 11 have been added to show where examples of these claim features are illustrated. Similar comments to those for claim 1 apply.

<u>Independent claim 21</u>

This claim specifies a communications network (figs 1a, 3a-3c, lines 17-25 of page 6) having a transport layer (fig 3b) and an overlay topology (fig 3a, lines 17-20 of page 8), in which adjacencies (12,14, lines 22-25 of page 8) are defined between a plurality of network nodes, wherein the network is provided with a path protection system for calculating recovery traffic paths so as to remedy network faults, said protection system being associated with a model (figs 3a-3c, lines 17-25 of page 8) of the network and being arranged to modify (102, fig 11, line 30 of page 14 onwards), within that model, the overlay topology by removal of selected adjacencies, to attempt computation (106, fig 11, lines 7-16 of page 15) of a path, and if (107, fig 11, lines 7-16 of page 15)no path is available, to remove fewer selected adjacencies (108, fig 11, lines 7-16 of page 15) from the overlay topology in said model and repeat said path computation (106).
Again, references to items shown in figures 1a, 2b, 3a, 3b, 3c and 11 have been added to show where examples of these claim features are illustrated. Similar comments to those for claim 1 apply.

(vi) Grounds of Rejection to be Reviewed on Appeal

There is one ground of rejection. Claims 1-28 have been rejected for obviousness over Jardetsky in view of Anderson. It seems that all rejections are dependent on whether Anderson shows the distinctive features of claim 1 or claim 3.

Even if the claim features or the disclosure of Anderson are somehow interpreted so that the answer is yes, the rejections are still then dependent on whether it would be obvious to combine the two references to reach the invention claimed.

(viii) Argument

Claim 1 and other claims are rejected for obviousness over a combination of Jardetsky and Anderson. The Examiner acknowledges that Jardetsky does not show the claim feature of removing fewer selected adjacencies from the overlay topology if no path is available, and repeating the path computation. Anderson is cited as showing these features, and it is alleged to have been obvious to incorporate these into Jardetsky.

As has been explained in the response of August 15, 2005, and not denied by the Examiner, Anderson and Jardetsky represent different approaches which are incompatible alternatives, so it would not be obvious to combine them. Even if they were combined, Anderson does not show all the missing features, and so the combination would not lead to the invention as claimed. The Examiner in the final office action of November 18, 2005 does not directly dispute these arguments, and can only make two indirect counter arguments at page 2 of the office action, as follows, neither of which are effective. Firstly, it is alleged that the applicants rely on features not claimed, and secondly it is alleged that Anderson shows the missing

features of removing fewer adjacencies from the overlay topology and repeating the path computation.

Regarding the first point, the features allegedly relied on but not claimed are obtaining weaker protection when there is no path available which avoids all the adjacencies. This is not an effective counter argument because as explained above in the summary of invention, obtaining weaker protection when there is no path available, is a result of the distinctive claimed feature of recovery path computation after "removing fewer selected adjacencies". The applicants do not rely on this result alone, the claim is distinguished by the above mentioned claim feature. The result of the claim feature is mentioned as it is relevant to obviousness, but nowhere have the applicants separated the result from the distinctive claim feature, so it is not clear how the Examiner can support an allegation that the applicants rely only on this result. To make this inference as to what the applicants rely on, either ignores the clear statement of the applicants that "neither Jardetsky nor Anderson show the claim feature of removing fewer adjacencies and retrying the path computation..." or assumes that the applicants are wrong about whether Anderson shows this distinctive claim feature.

This latter assumption relates to the second counter argument, which will now be addressed, the allegation that Anderson does show the missing features of removing fewer adjacencies from the overlay topology and repeating the recovery path computation. The Examiner cites paras 0028 and 0111, and furthermore para 0049 of Anderson, which read as follows:

> "[0028] FIG. 3B shows a representation of the forwarding table after the recovery path is activated by blockage of the primary path in accordance with an embodiment of the present invention. "
>
> "[0049] For example, with reference again to FIG. 1, Node A 102 preferably monitors the primary path 110 using a fast detection mechanism. Upon detecting a failure of the primary path 110, Node A 102 switches some or all of

the traffic from the primary path 110 to the recovery path 112. This typically involves inactivating the primary path 110 and activating the recovery path 112. This can be accomplished, for example, by removing the primary path from the forwarding table, blocking the primary path in the forwarding table, or marking the recovery path as a higher priority path than the primary path in the forwarding table."

"[0111] FIG. 8 shows exemplary logic 800 for performing a switch over from the primary path to the recovery path. Beginning at block 802, the logic may remove the primary path from the forwarding table, in block 804, block the primary path in the forwarding table, in block 806, or mark the recovery path as a higher priority path than the primary path in the forwarding table, in block 808, in order to inactivate the primary path and activate the recovery path. The logic may switch communications from the primary path to the recovery path based upon a predetermined priority scheme, in block 810. The logic 800 terminates in block 899."

The Examiner had earlier cited paras 0048 and 0078 of Anderson which were addressed in previous responses. None of these passages are concerned with computing the recovery path, they only concern inactivating the primary path and activating the recovery path. A skilled person would know that activating cannot be confused with nor "inferred" as a disclosure of computing the recovery path. Activating a recovery path can only occur once the recovery path has been computed. So these cited passages of Anderson cannot disclose computing a path.

Since the activating a recovery path can only take place if the path is available, these passages cannot show repeating a computation if no path is available. Claim 1 is explicitly limited to attempting computation of a path, and if no path is available, repeating the computation with a reduced number of adjacencies removed. This claim feature has a double negative, meaning the number of adjacencies involved in the computation, is enlarged for the repeat computation. The cited passages do not

involve computation of paths, but if path computation is somehow interpreted artificially as meaning looking up the pre computed path in a routing table, then there is still no disclosure of the claim feature of enlarging the number of adjacencies if no path is available.

In Anderson, the routing table would be altered so that a look up operation points either to the main path or the recovery path, there would not be a time when a result of a look up operation is "no path is available", and hence there is no disclosure of enlarging the number of adjacencies if no path is available.

Hence there cannot be a reasonable interpretation of the distinctive claim features which encompasses the disclosure of Anderson. Hence there is no support for the Examiner's counter argument that Anderson shows the features acknowledged as missing from Jardetsky. There is no other basis for an obviousness argument, as there is no incentive to alter Jardetsky to reach the present invention. Reference is made to previous responses for more detailed arguments.

To summarize some of the points made in previous responses, Anderson indicates that when a packet is unable to reach a destination, an alternate route is computed that does not share the next hop in a logical topology, with the original route, and this is repeated for calculating a next hop and so on, where there is a failure, until the destination is reached. In other words only the hops along the original logical path to that destination are removed when computing a backup path to use in the case of a failure. This is completely different from the method claimed and from the method of Jardetsky. Hence there is no incentive to select and combine parts of the known methods. In any case this would not result in the present invention.

The claim feature of retrying the path computation with fewer of the adjacencies removed, if necessary, can help enables a less optimal path to be found, when there is no optimal path. This may provide many grades of weaker protection and so

clearly makes better use of the available resources. There is no suggestion of this graded search for a successively weaker path, in Anderson or any of the cited documents.

Regarding independent claim 3, this claim recites a method of calculating a protection path by defining in a model of the network "a hierarchy of protection levels, each said protection level being characterized by a respective set of broken adjacencies in said model; attempting to calculate a recovery path for a selected protection level in said hierarchy; and if no said path is available, repeating said calculation attempt for successive further protection levels in said hierarchy until a protection path is identified."

The Examiner tries to argue that Anderson shows in paras 0048 and 0078 the claim features of if no path is available, repeating the calculation for further protection levels. There is mention of reporting failures to higher layers, but not of repeating a path calculation for further protection levels. This can bring similar advantages to those explained above in relation to claim 1. There is no suggestion in Anderson of this nor how to achieve the advantages. Hence this claim cannot be obvious over the cited documents taken alone or in combination.

Regarding independent claim 4, this has features corresponding to those of claim 3 and so is not anticipated for the same reasons.

Regarding independent claims 15, 16, 17, 18, 19 and 21, the same reasons as for claim 1 apply.

All the other claims are dependent or have corresponding distinctive features and so are all allowable for the same reasons. Reversal of the Examiner's rejections is respectfully requested.

November 13, 2006

Respectfully submitted,

William M. Lee

William M. Lee, Jr.
Registration No. 26,935
Barnes & Thornburg
P.O. Box 2786
Chicago, Illinois 60690-2786
(312) 214-4800
(312) 759-5646 (fax)

## Claims Appendix

1. A method of fault recovery in a multi-layer communications network having a transport layer topology and an overlay topology, in which adjacencies are defined between a plurality of network nodes, the method comprising, for each adjacency for which a recovery path is to be determined, modifying the overlay topology by removal of selected adjacencies, attempting computation of a path, and if no path is available removing fewer than the selected adjacencies from the overlay topology and repeating said path computation.

2. A method as claimed in claim 1, wherein said adjacencies are selected from a knowledge of the transport layer topology.

3. A method of calculating a protection path for traffic carried on a main path in a multilayer communications network having a lower transport layer and an upper layer incorporating a plurality of routers, and in which adjacencies are defined between respective pairs of routers, the method comprising the steps of;

defining a model of the network;

defining in said model a hierarchy of protection levels, each said protection level being characterised by a respective set of broken adjacencies in said model; attempting to calculate a recovery path for a selected protection level in said hierarchy; and

if no said path is available, repeating said calculation attempt for successive further protection levels in said hierarchy until a protection path is identified.

4. A method of calculating a protection path for traffic carried on a main path in a multi-layer communications network comprising a lower transport layer and an upper overlay incorporating a plurality of routers, there being a plurality of adjacencies defined between respective pairs of routers, wherein the method comprises the steps of;

defining a software model of the overlay of said network;

defining in said model a hierarchy of protection levels for said main path, each said protection level being characterised by a respective set of one or more broken adjacencies in said model;

selecting one said protection level and calculating a protection path avoiding the broken adjacencies associated with that protection level;

determining whether the calculated protection path is available in the network; and,

if said calculated path is not available in the network, repeating said path calculation and determining steps for one or more further selected protection levels.

5. A method as claimed in claim 4, wherein said protection levels are selected in order of hierarchy.

6. A method as claimed in claim 5, wherein said protection path is calculated via a next next hop algorithm.

7. A method as claimed in claim 6, wherein a protection level is selected according to a class of traffic carried on the path to be protected.

8. A method as claimed in claim 7, wherein said network incorporates a transport layer comprising a plurality of switches interconnected by optical fibre paths.

9. A method as claimed in claim 7, wherein the network model topology is defined by a first list of adjacencies representing the overlay topology, and a second list of paths, one for each adjacency.

10. A method as claimed in claim 9, and further comprising editing the network model topology by selecting sequentially the adjacencies in the overlay topology, testing each adjacency against assumptions about what equipment has failed in light of a hypothesised IP layer adjacency loss indication, and, if the adjacency passes the test, removing it from the topology.

11. A method as claimed in claim 10, wherein said transport layer comprises a synchronous network.

12. A method as claimed in claim 11, wherein said network is a packet network.

13. A method as claimed in claim 12, wherein said network is a multi-protocol label switched network.

14. A method as claimed in claim 13, wherein said network incorporates one or more virtual private networks.

15. A method of identifying a protection path for traffic carried on a main path in a multi-layer communications network so as to remedy a fault condition involving one or more network elements, the method comprising the steps of:

assuming the failure of a plurality of network elements;

calculating a protection path circumventing all the assumed failed network elements and determining whether that calculated path is an available path in the network; and

if said calculated path is not an available path in the network, successively assuming the failure of fewer than the assumed failed network elements and repeating the path calculation and determination steps for each successive assumption until an available protection path is identified.

16. Software in machine readable form for identifying a protection path for traffic carried on a main path in a multi-layer communications network so as to remedy a fault condition involving one or more network elements, the software being arranged to perform the method steps of:

assuming the failure of a plurality of network elements;

calculating a protection path circumventing all the assumed failed network elements and determining whether that calculated path is an available path in the network; and

if said calculated path is not an available path in the network, successively assuming the failure of fewer than the assumed failed network elements and repeating the path calculation and determination steps for each successive assumption until an available protection path is identified.

17. A method of network management for planning fault recovery paths in a multi-layer communications network having a transport layer and an overlay topology, and in which adjacencies are defined between a plurality of network nodes, the method comprising, providing a model of the network, and within said model for each recovery path to be determined, modifying the overlay topology by removal of selected adjacencies, attempting computation of a recovery path, and if no recovery path is available removing fewer selected adjacencies from the overlay topology and repeating said recovery path computation.

18. A method of network management for planning fault recovery paths in a multi-layer communications network so as to remedy a future fault condition involving one or more network elements on a traffic path, the method comprising the steps of:

providing a model of the network

assuming the failure of a plurality of network elements in said model;

calculating a protection path circumventing all the assumed failed network elements and determining whether that calculated path is an available path in the network; and

if said calculated path is not an available path in the network, successively assuming the failure of fewer than the assumed failed network elements and repeating the path calculation and determination steps for each successive assumption until an available protection path is identified.

19. A network manager for a multi-layer communications network and for planning fault recovery paths in said network so as to remedy a fault condition involving one or more network elements on a traffic path, the network manager being arranged to perform the method steps of:

providing a model of the network

assuming the failure of a plurality of network elements in said model;

calculating a protection path circumventing all the assumed failed network elements and determining whether that calculated path is an available path in the network; and

if said calculated path is not an available path in the network, successively assuming the failure of fewer than the assumed failed network elements and repeating the path calculation and determination steps for each successive assumption until an available protection path is identified.

20. A network manager as claimed in claim 19, and embodied as software in machine readable form on a storage medium.

21. A communications network having a transport layer and an overlay topology, in which adjacencies are defined between a plurality of network nodes, wherein the network is provided with a path protection system for calculating recovery traffic paths so as to remedy network faults, said protection system being associated with a model of the network and being arranged to modify, within that model, the overlay topology by removal of selected adjacencies, to attempt computation of a path, and if no path is available, to remove fewer than the selected adjacencies from the overlay topology in said model and repeat said path computation.

22. A network as claimed in claim 21, wherein said protection system defines in said model a hierarchy of protection levels for said main path, each said protection level being characterised by a respective set of one or more broken adjacencies in said model.

23. A network as claimed in claim 22, wherein said protection system selects said protection levels in order of hierarchy.

24. A network as claimed in claim 23, wherein said protection path is calculated via a next next hop algorithm.

25. A network as claimed in claim 24, and incorporating a transport layer comprising a plurality of switches interconnected by optical fibre paths.

26. A network as claimed in claim 25, wherein said transport layer comprises a synchronous network.

27. A network as claimed in claim 26, and comprising a packet network.

28. A network as claimed in claim 27, and comprising a multi-protocol label switched network.

**Evidence Appendix**

There is no such appendix.

## Related Proceeding Appendix

There is no such appendix.